

IN THE CLAIMS:

This listing of claims will replace all prior versions, and listings, of claims in the present application:

1. (Currently amended) A method for mirroring data comprising:

receiving at a first storage server a data access request from a client coupled to the first storage server;

writing the data access request to a first ~~portion of a non-volatile storage device log in a~~ first non-volatile random access memory in the first storage server, to prevent a loss of the data access request;

transmitting the data access request from the first storage server to a second storage server, wherein the second storage server writes the data access request to a second non-volatile log in a second non-volatile random access memory in the second storage server and into a file stored in a mass storage device managed by the second storage server; and

when the first ~~portion of the non-volatile storage device log~~ in the first storage server becomes full,

~~applying writing~~ the data access request stored in the first ~~portion of the non-volatile storage device log~~ to a first volume managed by the first storage server, and

causing the second storage server to ~~apply write~~ the data access request stored in the file ~~stored in the mass storage device to a second volume managed by the second storage server, wherein the second volume is an image volume of the first volume, wherein the second storage server manages the image volume and the mass storage device,~~ wherein the second storage server uses the file to recover data in the image volume if a disaster occurs.

2. (Canceled)

3. (Currently amended) The method of claim 1, further comprising:

 sending a synchronization request to the second storage server from the first storage server when the first ~~portion of the non-volatile storage device~~ log in the first storage server ~~[[is]]~~ becomes full.

4. (Previously presented) The method of claim 1, further comprising:

 sending an acknowledgement from the second storage server to the first storage server in response to receiving the data access request to cause the first storage server to send a response to the client after the data access request has been stored on the first storage server and stored in the mass storage device managed by the second storage server.

5. (Canceled)

6. (Currently amended) The method of claim 1, wherein the file is associated with the first ~~portion of the non-volatile storage device~~ log in the first storage server.

7. (Previously presented) The method of claim 1, wherein the data access request is transmitted from the first storage server to the second storage server over a network.

8. (Previously presented) The method of claim 1, further comprising:

assigning a sequence number to the data access request, wherein the sequence number designates a position of the data access request in a group of data access requests to ensure that the data access request is properly ordered within the file.

9. (Canceled)

10. (Currently amended) An apparatus comprising:

a destination storage server to mirror data stored by a source storage server;

a network interface in the destination storage server, the network interface coupled to the source storage server, the network interface to receive a data access request from a client coupled to the source storage server, wherein the source storage server has written the data access request ~~[[in]]~~to a first ~~portion of a non-volatile storage device~~log in a first non-volatile random access memory in the source storage server to prevent a loss of the data access request, wherein the destination storage server is configured to write the data access request to a second non-volatile log in a second non-volatile random access memory in the destination storage server and to a file corresponding to the source storage server; and

a mass storage device₁ managed by the destination storage server₁ to store the file, wherein the data access request is ~~applied-written~~ to ~~[[a]]the non-volatile-mass storage device managed by the destination storage server~~ in response to the first ~~portion of the non-volatile storage device~~log in the source storage server becoming full, wherein the destination storage server uses the file to recover data in the nonvolatile mass storage device when a disaster occurs.

11. (Canceled)

12. (Previously presented) The apparatus of claim 10, wherein the network comprises a Transmission Control Protocol/Internet Protocol (TCP/IP) network.

13. (Canceled)

14. (Canceled)

15. (Currently amended) The apparatus of claim 10, wherein the destination storage server modifies an image of a volume, said volume [[is]] maintained by the source storage server, according to the data access request when the source storage server makes a synchronization request, wherein the image of the volume is in the non-volatile mass storage device managed by the destination storage server.

16. (Canceled)

17. (Previously presented) A method comprising:

receiving a data access request at a destination filer from a first source filer, wherein the data access request is written to a first portion of a first nonvolatile random access memory in the source filer;

sending an acknowledgement to the first source filer in response to the destination filer receiving the data access request;

writing the data access request to a second nonvolatile random access memory in the destination filer;

transferring the data access request from the second nonvolatile random access memory to a first file, corresponding to the first source filer, on a mass storage device managed by the destination filer;

removing the data access request from the second nonvolatile random access memory after transferring the data access request to the first file;

applying the data access request in the first file to an image of a volume in response to a specified signal from the first source filer indicating that the first portion of the first nonvolatile random access memory is full, wherein the volume is maintained by the first source filer and the image is maintained by the destination filer, wherein the destination storage server uses the first file to recover data in the image of the volume if a disaster occurs;

receiving a second data access request from a second source filer, wherein the second data access request is written to a third nonvolatile random access memory in the second source filer;

sending a second acknowledgement to the second source filer in response to the destination filer receiving the second data access request;

writing the second data access request to the second nonvolatile random access memory in the destination filer;

transferring the second data access request from the second nonvolatile random access memory to a second file on the mass storage device managed by the destination filer, the second file associated with the second source filer; and

removing the second data access request from the second nonvolatile random access memory after transferring the second access request to the second file.

18. (Canceled)

19. (Previously presented) The method of claim 17, further comprising connecting the second source filer to the client in response to a system failure.

20. (Previously presented) The method of claim 17, further comprising:
allowing the client to access the image.

21. (Currently amended) A method of mirroring data, the method comprising:
operating a destination storage server to maintain a plurality of mirror volumes in a non-volatile mass storage subsystem, wherein each mirror volume mirrors a corresponding one of a plurality of source volumes maintained by a plurality of source storage servers that are coupled to communicate with the destination storage server;

receiving, at the destination storage server, write requests from the plurality of source storage servers, each said write request corresponding to a write request received by one of the plurality of source storage servers from a storage client for updating one of the plurality of source volumes;

operating the destination storage server to store the write requests temporarily prior to synchronizing the plurality of mirror volumes with the plurality of source volumes, including

storing a non-volatile log of the write requests received by the destination storage server from the plurality of source storage servers in a non-volatile random access memory in the destination storage server,

using the destination storage server to maintain a plurality of files in a non-volatile mass storage subsystem, each said file corresponding to a separate one of the plurality of source storage servers, and

storing each write request received by the destination storage server from a source storage server in one of said files, which corresponds to said source storage server; and
in response to receiving a specified signal from the source storage server, operating the destination storage server to synchronize the plurality of mirror volumes with the plurality of source volumes based on the write requests received from the plurality of source storage servers, wherein each of the source storage servers maintains a separate non-volatile log of write requests from storage clients in a partitioned non-volatile random access memory, and wherein the specified signal from the source storage server corresponds to a partition becoming full in the partitioned non-volatile random access memory in the source storage servers.

22. (Canceled)

23. (Previously presented) A method as recited in claim 21, wherein each partition of the partitioned non-volatile random access memory in each of the source storage servers corresponds to a distinct one of the plurality of files in the non-volatile mass storage subsystem.

24. (Currently amended) A method for mirroring data comprising:

receiving at a source storage server a data access request from a client coupled to the source storage server;

writing the data access request to a first ~~portion of a non-volatile storage device~~log in a first non-volatile random access memory in the source storage server to prevent a loss of the data access request;

transmitting the data access request from the source storage server to a destination storage server, wherein the destination storage server writes the data access request into a second non-volatile log in a second non-volatile random access memory in the destination storage server and a file stored in a mass storage device managed by the destination storage server; and

when the first ~~portion of the non-volatile storage device~~log in the source storage server becomes full,

transmitting the data access request in the first ~~portion of the non-volatile storage device~~log to a first volume managed by the source storage server, to cause the first volume to be updated according to the data access request, and

causing the destination storage server to transmit the data access request in the file ~~stored in the mass storage device~~ to a second volume managed by the destination storage server, wherein the second volume is an image volume of the first volume, to cause the image volume to be updated according to the data access request.

25. (Currently amended) The method of claim 24, wherein the file is associated with the first ~~portion of the non-volatile storage device~~log in the source storage server.

26. (Currently amended) The method of claim 24, wherein the mass storage device ~~includes~~
comprises a disk.

27. (Previously presented) The method of claim 24, further comprising:

assigning a sequence number to the data access request, wherein the sequence number designates a position of the data access request in a group of data access requests to ensure that the data access request is properly ordered within the file.